

Achieving 99.999% Uptime with SchoonerSQL: How Schooner delivers high availability

Synopsis

In this whitepaper, learn more about high availability strategy and business continuity planning requirements. 99.999% uptime (or “five nines”) will be defined, discussed and illustrated using a variety of use cases. Learn how you can use SchoonerSQL’s mix of highly-parallel synchronous and asynchronous replication coupled with immediate and automated failover across LAN, MAN or WAN to realize 99.999% availability from MySQL and its standard InnoDB storage engine.

Table of Contents

1. About High Availability	1
2. SchoonerSQL Architecture & Virtual IP Management	1
3. SchoonerSQL Replication Flexibility Delivers 99.999% uptime	2
a) 99.999% uptime inside a local synchronous cluster.....	2
b) 99.999% uptime inside a datacenter, with multiple clusters	5
Case 1: If a Master fails inside a synchronous cluster	5
Case 2: If an async Slave fails in one of the synchronous cluster groups	6
c) 99.999% uptime across datacenters	8
Case 1: Automated failover when the update Master fails inside a datacenter	8
Case 2: Automated failover if an async Slave (receiving point) fails inside a datacenter ...	9
4. SchoonerSQL: Hardware and Network Setup.....	11
5. Conclusion	12

1. About High Availability

Reliable access to rapidly increasing amounts of data drives customer-facing applications such as web sites, most business processes, and many products. The database is therefore an essential part of the infrastructure of most datacenters today. Organizations cannot afford the lost revenue, customers, and reputation that result from application downtime. Ensuring year-round 24/7 database availability is a major operational challenge for datacenter managers.

Five 9s — or 99.999% uptime — is a popular term to describe the gold standard of high availability (HA). With 99.999% uptime a service will be down for approximately 5.26 minutes per year or 25.9 seconds per month or 6.05 seconds per week. 99.999% uptime is the highest practical level of availability and uptime for most organizations. But achieving it is not easy with any database, especially with traditional MySQL offerings.

Availability %	Downtime per year	Downtime per month	Downtime per week
90% (1 nines)	36.5 days	72 hours	16.8 hours
99% (2 nines)	3.65 days	7.20 hours	1.68 hours
99.9% (3 nines)	8.76 hours	43.2 minutes	10.1 minutes
99.99% (4 nines)	52.56 minutes	4.32 minutes	1.01 minutes
99.999% (5 nines)	5.26 minutes	25.9 seconds	6.05 seconds

This white paper describes the architecture of SchoonerSQL™, an enterprise-grade OLTP database that delivers high availability while maintaining 100% compatibility with MySQL and its standard InnoDB storage engine. The paper then explains how to use the flexible synchronous and asynchronous replication modes in SchoonerSQL to get five nines in most common database use cases. The paper then describes the recommended HA hardware setup for using SchoonerSQL.

2. SchoonerSQL Architecture & Virtual IP Management

SchoonerSQL is architected to provide 99.999% availability with great performance, effective scalability and easy management. Yet SchoonerSQL is 100% compatible with MySQL and its standard InnoDB storage engine, since it is built using the commercial MySQL and InnoDB source code.

The key elements of the SchoonerSQL architecture are:

- Highly-parallel synchronous replication across a LAN;
- Highly-parallel asynchronous replication across LAN, MAN or WAN;
- Immediate automated failover and recovery across LAN, MAN or WAN
- Use of Virtual IPs (VIPs) to provide service continuity during automated failover of your database servers;


SchoonerSQL eliminates the need to worry about manual failover or manual promotion of a surviving node to become the new master: our solution automates this process. Everything is easily managed and monitored through our Admin Cluster UI (available both in an elegant GUI and a full command line interface). SchoonerSQL also includes integrated online backup.

Replication for redundancy is typically accomplished using a Master - Slave configuration of servers. As with traditional MySQL, a SchoonerSQL cluster will have a Master server and multiple Slaves. But the traditional concept of a Slave is different in a SchoonerSQL synchronous cluster. The term Slave has been used in traditional MySQL because the slave servers have to perform every task in copying from the Master binary log, then updating their relay logs and finally committing to slave databases. The Master server plays no role in replication other than storing the replication events in the binary log.

SchoonerSQL instead does synchronous replication using a push mechanism during transaction execution. SchoonerSQL implements this using multiple parallel threads, which apply the replication events simultaneously on all of the Slaves. In this way every node in the cluster is 100% in-synch at all times. Thus it is more accurate to use the term “Read Masters” instead of Slaves when using SchoonerSQL, because reads from SchoonerSQL slave nodes are always up-to-date.

The SchoonerSQL approach has several benefits:

- **Read consistency across all nodes:** The Read Masters (aka Slaves) are always current with the Master so there's no stale data and no slave lag.
- **Zero data loss:** Even if a Master fails, the same data is immediately available on all other nodes in the cluster, so there is zero data loss.
- **Maximal performance:** There's no need to throttle back the Master to allow slower Slaves to keep up, resulting in vastly greater throughput with SchoonerSQL than with any other MySQL offering.




The diagram above shows that a Read Master's data is 100% consistent with the Master that receives and handles the writes, eliminating the possibility of stale data. We also use checksums to prevent data corruption and automate the failover, guaranteeing no data loss. Even if a Master receiving the writes fails, automatic failover happens and a Read Master will act as the actual Master receiving all the writes.

The administrator uses the SchoonerSQL Cluster Admin GUI to initially set the Write and Read VIPs (virtual IP addresses) to the appropriate servers. SchoonerSQL then handles failover automatically by transparently switching the VIPs. You can have many Read VIPs but only one write VIP in a given setup. Our VIP management automates the master failover process with full service continuity as seen in the next diagram. It shows a 3-node cluster where node 1 acts as Master and nodes 2 and 3 act as Read Masters.

Consider this failover scenario. When the Master in node 1 fails, the write VIP of the failed Master is automatically assigned to node 2 and a read VIP is assigned to node 3. The failover time is from 1 to 3 seconds and is totally automatic.

SchoonerSQL VIP Auto Failover



3. SchoonerSQL Replication Flexibility Delivers 99.999% Uptime

SchoonerSQL offers a robust solution to provide users with full service continuity. The heart of our solution is the full synchronous parallel replication explained earlier. SchoonerSQL performs highly-parallel replication and automated failover to provide five nines not only inside a synchronous cluster but also in asynchronous clusters inside or between datacenters. This makes users and customers happy, and makes administrators happy by eliminating the headaches of error-prone manual failover processes of stopping a slave, resetting it to master and so on.


There are multiple ways to use SchoonerSQL replication mechanisms to deliver 99.999% uptime:

a) 99.999% uptime inside a local synchronous cluster


SchoonerSQL supports up to 8 synchronous nodes in a cluster with one Master and seven Read Masters. Since SchoonerSQL keeps the data fully consistent across all eight nodes, even if the master fails, the write VIP will automatically failover to any of the 7 read masters within 1 to 3 seconds without service interruption. If a read master fails, its read VIPs are automatically switched to another read master server within 1 to 3 seconds.

The following diagram, using a 4-node synchronous cluster, shows how the master failover happens inside the local synchronous cluster.


Normal operation



Master goes down



New Master takes over within
1 to 3 seconds providing 99.999% uptime



Because the data is consistent across all the nodes all the time, any node can take over as the Master and start receiving writes. Priority and setting up of failover or fail back can be configured in the Admin Cluster UI.


b) 99.999% uptime inside a datacenter, with multiple clusters

To get replication flexibility and unlimited read scalability inside a datacenter, you simply use two or more SchoonerSQL synchronous clusters (up to 8 nodes in a cluster). Automatic failover, automatic node recovery, and automatic rejoining of failed nodes inside a synchronous cluster are tremendous benefits with such deployments.


Replication between multiple synchronous clusters inside a data center is based on SchoonerSQL parallel asynchronous replication, denoted as “async” in the diagrams below.

Case 1: If a Master fails inside a synchronous cluster


Normal operation:



Failure: The Master handling writes fails, and async replication breaks across the other clusters:




Resolution: SchoonerSQL ensures that a new Master automatically takes over within seconds and starts handling writes as seen below. The other synchronous clusters automatically join this new master within seconds, providing 99.999% availability:




Case 2: If an async Slave fails in one of the synchronous cluster groups:


Normal operation:



Failure Event: an async slave in the synchronous cluster fails, causing asynchronous replication between clusters to break:



Resolution: The new async slave takes over and rejoins the master within seconds, providing 99.999% availability:




c) 99.999% uptime across datacenters


The much higher network latencies of WAN compared to a LAN make it impractical to use synchronous replication across datacenters; the application would slow down too much. However, SchoonerSQL's unique use of parallel appliers across WAN increases the throughput tremendously so that the async slaves in the other datacenter receive updates with the minimum practical lag. SchoonerSQL also provides automated failover across WAN or across data centers and eliminates the downtime and hassle of doing a manual failover.

Case 1: Automated failover when the update Master fails inside a datacenter


Normal Operation:



Event: Master receiving writes fails and asynchronous replication breaks across DCs:




Resolution: SchoonerSQL cluster management ensures that the async slave across datacenter 2 automatically connects to this new master within seconds, providing 99.999% availability:




Case 2: Automated failover if an async Slave (receiving point) fails inside a datacenter


Normal Operation:



Event: An async slave in datacenter 2 fails, breaking asynch replication between datacenters:




Resolution: SchoonerSQL ensures that an existing node in datacenter 2 automatically takes over within seconds and acts as a receiving point. The new async slave across datacenter 2 automatically connects to the master in datacenter 1 within seconds as shown below, thus providing 99.999% availability:



4. SchoonerSQL: Hardware and Network Setup

SchoonerSQL is the most reliable high-availability implementation of MySQL and InnoDB on the market. But to get 99.999% availability, you need redundancy in your hardware too.



SchoonerSQL customers use these recommended best practices to get the full benefit of SchoonerSQL:

- Redundant network links, network switches, and power circuits
 - Eliminate single points of HW failure
- Distinct client and database replication network interfaces
 - Ensure sufficient bandwidth capacity for replication
 - Avoid cascaded or dependency-based failures
- Network bonds and VIPs
 - Seamless and instantaneous failover
- Serial / USB / secondary Ethernet connection
 - Simple method for breaking a connection, when needed
- SchoonerSQL achieves system-level redundancy and failover without data loss
 - Detects system failure
 - Maintains health information on all configured network interfaces
 - Detects split brain
 - Manages Virtual IPs
 - Provides policy-driven instantaneous failover

5. Conclusion

SchoonerSQL offers an ideal solution for your application if you need 99.999% availability and want the migration convenience of 100% compatibility with MySQL and InnoDB (requiring no application rewrites). SchoonerSQL has the necessary features to keep your service or website up and running with just minutes of downtime per year. Schooner offers a truly mission-critical solution that is broadly deployed across all industry segments, with customers realizing the business benefits 24x7x365.

Here is a simple summary of SchoonerSQL's advantages in mission-critical MySQL:

99.999% High Availability

- Slash unplanned downtime with immediate automatic fail over
- Slash planned downtime with automated upgrade and migration
- Extremely flexible replication modes

Hassle-Free Operations

- Immediate automated failover, recovery across LAN, MAN, or WAN
- No error-prone manual processes for failover, recovery, provisioning

Great Performance

- Get the most out of your HDDs, SSDs, or SAN
- 4x more throughput vs. MySQL 5.5
- Reduce your server footprint
- High performance WAN Support

Highest Data Integrity

- No slave lag
- No data loss
- No stale data
- No corrupt data

Highest Scalability

- Easy vertical scaling using processor cores and SSDs
- Unlimited linear read scaling
- Unlimited write scaling (dbShards)

Easiest Management

- 1-click node addition, removal, promotion
- Extensive monitoring capabilities
- Easy alerts
- Integrated online backup

100% MySQL and InnoDB Compatible

Evaluating the options and trade-offs for your Data Center?

Keep in Touch:

info@schoonerinfotech.com

www.schoonerinfotech.com

Tel: +1 408-773-7500

www.schoonerinfotech.com

About Schooner Information Technology

Schooner Information Technology delivers cost-effective, enterprise-grade database software that frees customers from the traditional tradeoffs between availability, performance and total cost of ownership (TCO). Schooner's proven, cost-effective offerings include SchoonerSQL, which is 100% compatible with MySQL for easy adoption, and the complementary Membrain data store and software cache. Based in Sunnyvale, California, Schooner is privately funded and supports a variety of customers across technology, entertainment, financial services and telecommunications industries.

For more information, please visit www.schoonerinfotech.com.

Follow us on Twitter @schoonerinfo.

Schooner Information Technology, Inc.

501 Macara Ave., Suite 101
Sunnyvale, CA 94085

Main: (408) 773-7500

Service and Support: (877) 888-5064

Fax: (408) 736-4212